arXiv:1702.05139v1 [cs.CC] 16 Feb 2017

On the Bit Complexity of Sum-of-Squares Proofs

Prasad Raghavendra* Ben Weitz†

February 20, 2017

Abstract

It has often been claimed in recent papers that one can find a degree d Sum-of-Squares proof if one exists via the Ellipsoid algorithm. In [O17], Ryan O'Donnell notes this widely quoted claim is not necessarily true. He presents an example of a polynomial system with bounded coefficients that admits low-degree proofs of non-negativity, but these proofs necessarily involve numbers with an exponential number of bits, causing the Ellipsoid algorithm to take exponential time. In this paper we obtain both positive and negative results on the bit complexity of SoS proofs.

First, we propose a sufficient condition on a polynomial system that implies a bound on the coefficients in an SoS proof. We demonstrate that this sufficient condition is applicable for common use-cases of the SoS algorithm, such as Max-CSP, Balanced Separator, Max-Clique, Max-Bisection, and Unit-Vector constraints.

On the negative side, O'Donnell asked whether every polynomial system containing Boolean constraints admits proofs of polynomial bit complexity. We answer this question in the negative, giving a counterexample system and non-negative polynomial which has degree two SoS proofs, but no SoS proof with small coefficients until degree $\Omega(\sqrt{n})$.

*UC Berkeley, raghavendra@berkeley.edu. Supported by NSF Career Award, NSF CCF-1407779 and the Alfred. P. Sloan Fellowship.

†UC Berkeley, bsweitz@cs.berkeley.edu. Supported by an NSF Graduate Research Fellowship (NSF award no DGE 1106400).

1 Introduction

The Sum of squares (SoS) proof system is a versatile and powerful approach to certifying polynomial inequalities. SoS certificates can be shown to underly a vast number of algorithms in combinatorial optimization. On the one hand, SoS certificates hold the promise of yielding algorithms that possibly refute the notorious unique games conjecture [BBH$^+$12, BRS11, GS11]. On the other hand, a flurry of recent works have applied SoS proofs to develop algorithms for problems ranging from constraint satisfaction problems to tensor problems.

To illustrate sum of squares certificates, let us consider the example of the Balanced Separator problem. Here we are given a graph $G = (V, E)$ and the goal is to find a balanced cut $(S, \overline{S})$ with the minimum number of crossing edges. Like many problems in combinatorial optimization, it can be reformulated as a low-degree polynomial optimization problem. Specifically if we associate $\{0, 1\}$ variables $\{x_1, \ldots, x_n\}$ for the vertices of the graph G then we can rewrite the Balanced Separator problem as follows,

$$\text{Minimize} \quad \sum_{(i,j)\in E} (x_i - x_j)^2 \qquad \text{subject to } \left\{ x_i^2 = x_i \forall i \ , \ \frac{n}{3} \leq \sum_i x_i \leq \frac{2n}{3} \right\}$$

Here the constraint $x_i^2 = x_i$ ensures $x_i \in \{0, 1\}$ while the inequalities enforce the condition that the cut is balanced. More generally, a low-degree polynomial optimization is of the form

$$\begin{aligned} &\text{Minimize } r(x) \text{ subject to} \\ &\text{equalities } \mathcal{P} = \{p_i(x) = 0 | i \in [n]\} \text{ and inequalities } \mathcal{Q} = \{q_i(x) \geq 0 | i \in [m]\} \end{aligned}$$

An SoS certificate of a lower bound $r(x) \geq \theta$ is given by a polynomial identity of the form

$$r(x) - \theta = \sum_i h_i(x)^2 + \sum_{i\in[m]} \left(\sum_j^{t_i} s_j^2(x) \right) \cdot q_i(x) + \sum_{i\in[n]} \lambda_i(x) p_i(x)$$

Notice that for all x satisfying the equalities $\mathcal{P}$ and the inequalities $\mathcal{Q}$, the right hand side of the above identity is manifestly non-negative, thereby certifying that $r(x) \geq \theta$. The degree of the SoS certificate is the maximum degree of the polynomials involved, i.e., $d = \max\{\deg h_i^2, \deg s_j^2 q_i, \deg \lambda_i p_i\}$.

The main appeal of SoS certificates for polynomial optimization is that the existence of a degree d SoS certificate can be formulated as the feasibility of a semidefinite program (SDP). This is the degree d SoS relaxation first introduced by Shor [Sho87], and expanded upon by later works of Nesterov [Nes00], Grigoriev and Vorobjov [GV01], Lasserre [Las00, Las01] and Parrilo [Par00]. (see, e.g., [Lau09, BS14] for many more details).

The degree d SoS SDP has $n^{O(d)}$ variables, and if the coefficients of p and q are reasonably bounded (smaller than $2^{n^{O(d)}}$), the resulting SDP has a compact description of size $n^{O(d)}$. From this, several works including those by the authors, asserted that the resulting feasibility SDP can be solved in time $n^{O(d)}$ using the Ellipsoid algorithm.

In a recent work, O'Donnell [O17] observed that this often repeated claim is far from true. Specifically, O'Donnell exhibited systems of polynomial inequalities with bounded coefficients such that only degree 2 SoS certificates of non-negativity involve coefficients that are doubly exponential in size. Thus all SoS certificates need an exponential number of bits to represent and consequently, the ellipsoid algorithm will incur an exponential running time.

As pointed out by O'Donnell, the issue at hand here is not just that of additive error in the solution, i.e., the difference between testing feasibility and near-feasibility. Indeed, semidefinite programming via the ellipsoid algorithm can only test feasibility up to a very small additive error. However, in a majority of applications of SoS SDP relaxations in combinatorial optimization, the variables in the underlying polynomial system are explictly bounded (also known as Archimedian). Specifically, these include constraints such as $\{x_i^2 \leq 1 | i \leq [n]\}$, which yield explicit bounds on the values of the variables. In these settings, if there is an approximate SoS certificate for $r(x) \geq \theta$, then there exists a proper SoS certificate for a slightly weaker lower bound $r(x) \geq \theta - o(1)$. Therefore, additive error incurred in semidefinite programming can often be traded off for a slightly weaker objective value. The issue highlighted by O'Donnell is far more serious in that the coefficients of the SoS certificate are too large – thereby directly affecting the runtime of the ellipsoid algorithm.

On a positive note, O'Donnell shows that a polynomial system whose only constraints are the Boolean constraints $\{x_i^2 = x_i | i \in [n]\}$ always admit SoS certificates with polynomial bit complexity. He proceeds to ask whether all polynomial systems that include boolean constraints, potentially among others, always admit bounded SoS certificates.

1.1 Our Results

In this work, we further explore the issue of bit complexity of SoS proofs, and obtain both positive and negative results.

First, we present an easily verifiable and broadly applicable set of sufficient conditions under which a polynomial optimization problem has small SoS certificates. Roughly speaking, we show that polynomial systems with *rich* sets of solutions have bounded SoS certificates of non-negativity. Consider a system consisting of polynomial equalities $\mathcal{P}$ and inequalities $\mathcal{Q}$. Our approach consists of looking for assignments S satisfying three criteria (see Definition 2.3 and Theorem 4.1 for formal statements).

Theorem 1.1. *Assume $(\mathcal{P}, \mathcal{Q}, S)$ satisfies:*

1. *The assignments S robustly satisfy the inequalities in $\mathcal{Q}$.*
2. *The polynomial calculus proof system is both complete and efficient over S. In other words, all degree d polynomial identities over S can be derived using a degree $O(d)$ polynomial derivation from the equalities $\mathcal{P}$.*
3. *The assignments S are spectrally rich in that smallest non-zero eigenvalue of their covariance matrix is at least $2^{-\operatorname{poly}(n^d)}$.*

Then if r has a degree d proof of non-negativity from $\mathcal{P}$ and $\mathcal{Q}$, it also has a degree $O(d)$ proof of non-negativity with coefficients bounded by $2^{\operatorname{poly}(n^d)}$.

We demonstrate the broad applicability of the above set of sufficient conditions by using them to show upper bounds on bit complexity for Max-CSP, Max-Clique, Matching, Balanced Separator, Max-Bisection, and optimization over the unit sphere. In each case, the above sufficient conditions can be verified easily.

The above set of sufficient conditions are widely applicable in combinatorial optimization, wherein the polynomial system is typically a relaxation of a well-known set of integer solutions. In such a setup with integer solutions, we observe in Section 3 that spectral richness is an immediate consequence of the discrete nature of the set of solutions. Therefore, in all these setups, the only non-trivial thing to verify is the efficiency of the polynomial calculus proof system.

The work of O'Donnell [O17] exhibited a polynomial system with bounded coefficients which admitted degree 2 SoS certificate, whose coefficients were necessarily doubly-exponential. However, the variables in this polynomial system were not all boolean, i.e. did not have the $x_i^2 = x_i$ constraint. In fact, O'Donnell asked whether every polynomial system with boolean constraints admits a small SoS proof. Moreover, the polynomial system in [O17] admits a degree 4 SoS certificate with small bit complexity. This opens up the possibility that one can effectively reduce the bit-complexity by raising the degree of the proof. For instance, if a system admits a degree d SoS certificate then does it always admit a degree 2^d SoS certificate with small bit complexity (even under boolean constraints)? Unfortunately, we refute both of the above possibilities by exhibiting a counterexample. Formally, we show the following:

Theorem 1.2. *There exists a system of quadratic equations on n variables such that*

- *The system includes the equation $x_i^2 - x_i = 0$ for each $i \in [n]$.*
- *There exists a polynomial with a degree 2 SoS certificate of non-negativity, albeit with doubly exponentially large coefficients.*
- *No SoS certificate of degree $d \leq \sqrt{n}$ has coefficients smaller than $\Omega\left(\frac{1}{n^d} \cdot 2^{\exp(\sqrt{n})}\right)$.*

2 Preliminaries

For a set of real polynomials $\mathcal{P} = \{p_1, p_2, \ldots, p_m\}$, we denote their generated ideal in $\mathbb{R}[x]$ by $\langle \mathcal{P} \rangle$ or $\langle p_1, \ldots, p_m \rangle$. We will be working with systems of polynomial constraints, and we will use the $\mathcal{P}$ to denote the equality constraints, and $\mathcal{Q}$ to denote the inequality constraints, i.e. $p(x) = 0$ and $q(x) \geq 0$ for $p \in \mathcal{P}$ and $q \in \mathcal{Q}$. We will usually use S for the set of points satisfying these constraints. We use $\mathbb{R}[x]_d$ for the set of polynomials of degree at most d, and we write $\mathbf{v}_d$ for the vector of polynomials whose entries are the elements of the usual monomial basis of $\mathbb{R}[x]_d$. Similarly, we use $\mathbf{v}(\alpha)_d$ for the vector of reals whose entries are the entries of $\mathbf{v}$ evaluated at α. We usually omit the dependencies on d as it is clear from context.

2.1 Polynomial Proofs

Let $\mathcal{P} = \{p_1, \ldots, p_n\}$ be a set of polynomials, and let $S = \{x \in \mathbb{R}^n | \forall p \in \mathcal{P} : p(x) = 0\}$. We define a proof of membership in $\langle \mathcal{P} \rangle$ as follows:

Definition 2.1. We say that $r(x)$ has a *derivation* from $\mathcal{P}$ if there is a polynomial identity of the form

$$r(x) = \sum_i^n \lambda_i(x) p_i(x).$$

We say that the proof has degree d if $\max_i \{\deg \lambda_i p_i\} = d$.

The above proof system is useful for proving when polynomials are zero on S, but often we want to prove that they are positive. To that end, let $\mathcal{P} = \{p_1, \ldots, p_n\}$ and $\mathcal{Q} = \{q_1, \ldots, q_m\}$ be two sets of polynomials, and let $S = \{x \in \mathbb{R}^n | \forall p \in \mathcal{P} : p(x) = 0, \forall q \in \mathcal{Q} : q(x) \geq 0\}$. We define a proof of non-negativity as follows:

Definition 2.2. We say that $r(x)$ has a *Sum-of-Squares proof of non-negativity from $\mathcal{P}$ and $\mathcal{Q}$* if there is a polynomial identity of the form

$$r(x) = \sum_i^{t_0} h_i^2(x) + \sum_i^m \left(\sum_j^{t_i} s_j^2(x) \right) q_i(x) + \sum_i^n \lambda_i(x) p_i(x).$$

We say the proof has degree d if $\max\{\deg h_i^2, \deg s_j^2 q_i, \deg \lambda_i p\} = d$.

The idea behind this terminology is that if such a proof exists, then r must be non-negative on S since the first two terms are non-negative, and the last term is zero. We will be concerned with not just the degree of these proofs, but also their bit complexity. To this end, we define the following norms on polynomials and proofs: For $p(x) \in \mathbb{R}[x]$, we write $\|p\|$ for the absolute value of the maximum coefficient of p in the standard monomial basis, and for any collection of polynomials $\mathcal{P}$, we write $\|\mathcal{P}\| = \max_{p \in \mathcal{P}} \|p\|$. We will also require a bound on $\|S\| = \max_{\alpha \in S} \|\alpha\|$. Throughout this paper we will assume that the solutions α are *explicitly bounded* by $\|\alpha\| \leq 2^{\text{poly}(n^d)}$.

2.2 Rich Solution Spaces

In this section we define the conditions we require in order to guarantee that SoS proofs from $\mathcal{P}$ and $\mathcal{Q}$ have low bit-complexity. For a set of assignments S to a polynomial system $(\mathcal{P}, \mathcal{Q})$, define the moment matrix as

$$M_S \coloneqq E_{\alpha \in S}[\mathbf{v}(\alpha)\mathbf{v}(\alpha)^T],$$

where the expectation is over the uniform distribution over S. We will omit the subscript and write M, if S is clear from the context.

Definition 2.3. With the above definitions,

- We say that S is *δ-spectrally rich for $(\mathcal{P}, \mathcal{Q})$ up to degree d* if every nonzero eigenvalue of M_S is at least δ.
- We say that $(\mathcal{P}, \mathcal{Q})$ is *k-complete on S up to degree d* if every zero eigenvector c of M_S (which can be seen as a degree d polynomial $c^T\mathbf{v}$) has a degree k derivation from $\mathcal{P}$.
- We say that S is *ε-robust for $\mathcal{Q}$* if $\forall q \in \mathcal{Q}, \forall \alpha \in S : q(\alpha) > \varepsilon$.

Spectral richness of the solutions S is equivalent to requiring if $p(x)$ is small on S, then there is a polynomial q which agrees with p on S and that has small coefficients. If $(\mathcal{P}, \mathcal{Q}, S)$ satisfies all three conditions then we say that S is *(δ, k, ε)-rich* for $(\mathcal{P}, \mathcal{Q})$ up to degree d. If $1/\delta = 2^{\text{poly}(n^d)}$, $k = O(d)$, and $1/\varepsilon = 2^{\text{poly}(n^d)}$ we simply say S is rich for $(\mathcal{P}, \mathcal{Q})$. We choose these bounds because Theorem 4.1 will imply that any constraints with a rich solution space has proofs of non-negativity that can be taken to have polynomial bit complexity. Before we get into the proof of the main theorem, we exhibit polynomial systems that admit rich solutions.

3 Examples with Rich Solution Spaces

In this section we present examples of polynomial systems that admit rich solution spaces. First, we consider the case $S \subseteq \{0,1\}^n$. In this case, the spectral richness is a consequence of the following easy observation.

Lemma 3.1. *Let $M \in \mathbb{R}^{N\times N}$ be an integer matrix with $|M_{ij}| \le B$ for all $i, j \in [N]$. The smallest non-zero eigenvalue of M is at least $(BN)^{-N}$.*

Proof. Let A be a full-rank principal minor of M and w.l.o.g. let it be at the upper-left block of M. We claim the least eigenvalue of A lower bounds the least nonzero eigenvalue of M. Since M is symmetric, there must be a C such that

$$M = \begin{bmatrix} I \\ C \end{bmatrix} A \begin{bmatrix} I & C^T \end{bmatrix}.$$

Let $P = [I, C^T]$, ρ be the least eigenvalue of A, and x be a vector perpendicular to the zero eigenspace of A. Then we have $x^T M x \ge \rho x^T P^T P x$, but x is also perpendicular to the zero eigenspace of $P^T P$. Now $P^T P$ has the same nonzero eigenvalues as $PP^T = I + C^T C \succeq I$, and thus $x^T P^T P x \ge 1$, and so every nonzero eigenvalue of A is at least ρ. Now A is a full-rank bounded integer matrix with dimension at most N. The magnitude of its determinant is at least 1 and all eigenvalues are at most $N \cdot B$. Therefore, its least eigenvalue must be at least $(BN)^{-N}$ in magnitude. □

Lemma 3.2. *Let $\mathcal{P}$ and $\mathcal{Q}$ be such that $S \subseteq \{0,1\}^n$. Then S is δ-spectrally rich with $\frac{1}{\delta} = 2^{\mathrm{poly}(n^d)}$.*

Proof. Recall $M = E_{\alpha\in S}[\mathbf{v}(\alpha)\mathbf{v}(\alpha)^T]$, and note that $|S| \cdot M$ is an integer matrix with entries at most 2^n. The proof follows by applying Lemma 3.1. □

To prove completeness, we typically want to show two things. First, that every degree d polynomial in $\langle \mathcal{P} \rangle$ has a degree at most k derivation. Second, that there are no polynomials outside $\langle \mathcal{P} \rangle$ that are zero on S. This second condition can be thought of as saying that the set of equations $\mathcal{P}$ is somehow maximal, i.e., if there are extra polynomial equalities implied by $\mathcal{Q}$, they should be included in $\mathcal{P}$. Here we consider a few examples.

Max-CSP: $\mathcal{P} = \{x_i^2 - x_i | i \in [n]\}$**.** Here $S = \{0,1\}^n$. Any polynomial p of degree d can be multilinearized one monomial at a time. Specifically, we can find degree d multilinear p^* such that $p - p^* = 0$ has a degree d derivation from $\mathcal{P}$. Furthermore, the multilinear polynomial p^* is zero over S if and only if all its coefficients are zero. Thus $\mathcal{P}$ is d-complete up to degree d for all $d \in \mathbb{N}$.

Max-Clique: $\mathcal{P} = \{x_i^2 - x_i | i \in [n]\} \cup \{x_i x_j | (i,j) \notin E\}$**.** Here S is the set of all cliques in the graph. Suppose p is a polynomial that is identically zero over S. Without loss of generality, we may assume p is multilinear, if otherwise we can multilinearize it using $\{x_i^2 - x_i | i \in [n]\}$. For a multilinear polynomial $p(x) = \sum_{\alpha\subset[n]} \hat{p}_\alpha x_\alpha$, we claim that if $p(x) = 0 \forall x \in S$ then for all cliques $\alpha \subset [n]$, the corresponding coefficient $\hat{p}_\alpha = 0$, i.e., all non-zero coefficients of p are non-cliques. Suppose not, then let α be the smallest clique with $\hat{p}_\alpha \neq 0$. Then, we will have $p(\mathbb{I}_\alpha) = \hat{p}_\alpha \neq 0$, a contradiction. Since all coefficients of p are non-cliques, each monomial in p can be eliminated using an appropriate polynomial from $\{x_i x_j | (i,j) \notin E\}$.

Remark 3.3. More generally, the above two cases are special cases of the following general setup: $\mathcal{Q}$ is empty, and $\mathcal{P}$ is a Gröbner basis. A Gröbner basis for an ideal is a generating set of polynomials that allow a well-defined multivariate polynomial division (see [AL94] for more information). Computing the Gröbner basis is often the first step in practical polynomial equation solvers, and we note the following easy lemma:

Lemma 3.4. *If $\mathcal{Q} = \emptyset$ and $\mathcal{P}$ is a Gröbner basis for $\langle \mathcal{P} \rangle$, then S is d-complete up to degree d.*

Proof. If $\mathcal{P}$ is a Gröbner basis, then every degree d polynomial in $\langle \mathcal{P} \rangle$ has a degree d derivation via multivariate division. Because $\mathcal{Q} = \emptyset$, the polynomials that are zero on S are exactly the polynomials in $\langle \mathcal{P} \rangle$. □

Balanced Separator: $\mathcal{P} = \{x_i^2 - x_i | i \in [n]\}$, $\mathcal{Q} = \{2n/3 - \sum_i x_i, \sum_i x_i - n/3\}$. The solution space S here is all bit strings with hamming weight between $n/3$ and $2n/3$. Suppose r is a polynomial that is zero on S. Without loss of generality, we may assume that r is multilinear by using the constraints $\{x_i^2 - x_i | i \in [n]\}$. Suppose r is a non-zero multilinear polynomial which is zero on S, then its symmetrized version $r^* = \frac{1}{n!} \sum_{\sigma \in \mathcal{S}_n} \sigma r$ must also be zero on S, where σ acts by permuting the variable names. However, r^* is a univariate polynomial in $\sum_i x_i$ (modulo the Boolean constraints). This univariate polynomial has $n/3$ zeros, and thus must have degree at least $n/3$. Since symmetrizing doesn't change degree, we conclude that r also has degree at least $n/3$. Thus every non-zero multilinear polynomial that is zero on S but not in $\langle \mathcal{P} \rangle$, has degree at least $n/3$. Therefore the system is d-complete up to degree d for $d \leq \frac{n}{3}$. The polynomials in $\mathcal{Q}$ can be perturbed by $1/2$ to make them $1/2$-robust, and thus S is rich for $(\mathcal{P}, \mathcal{Q})$.

Matching: $\mathcal{P} = \{x_{ij}^2 - x_{ij} | i, j \in [n]\} \cup \{\sum_i x_{ij} - 1 | i \in [n]\} \cup \{x_{ij} x_{ik} | i, j, k \in [n]\}$. These constraints are $2d$-complete as proven in [BBCH+16].

Max-Bisection: $\mathcal{P} = \{x_i^2 - x_i | i \in [n]\} \cup \{\sum_i x_i - n/2\}$. We will prove in Section 6 that these constraints are d-complete. The proof will be very similar to the one for MATCHING, due to the similar symmetry of the constraints.

Unit-Vector: $\mathcal{P} = \{\sum_i x_i^2 - 1\}$. Here $S = \{x : \|x\| = 1\}$. This constraint appears frequently in tensor norm problems as a way to enforce scaling. Since $\mathcal{Q} = \emptyset$, it is clearly robust. It may be well-known that $\mathcal{P}$ is d-complete, but we could not find a reference so we record it here for completeness. Let $p(x)$ be any degree d polynomial which is zero on the unit sphere, and define $p_0(x) = p(x) + p(-x)$. Clearly p_0 is also zero on the unit sphere, with degree $k = 2\lfloor (d+1)/2 \rfloor$. Note that p_0 has only terms of even degree. Define a sequence of polynomials $\{p_i\}_{i \in \{0,\ldots,k\}}$ as follows. Define q_i to be the part of p_i which has degree strictly less than k, and let $p_{i+1} = p_i + q_i \cdot (\sum_i x_i^2 - 1)$. Then each p_i is zero on the unit sphere and has no monomials of degree strictly less than $2i$. Thus $p_{k/2}$ is homogeneous of degree k. But then $p(tx) = t^k p_k(x) = 0$ for any unit vector x and $t > 0$, and thus $p_k(x)$ must be the zero polynomial. This implies that p_0 is a multiple of $\sum_i x_i^2 - 1$. The same logic shows that $p(x) - p(-x)$ is also a multiple of $\sum_i x_i^2 - 1$, and thus so is $p(x)$. Now $\langle \mathcal{P} \rangle$ is principal, so every degree d polynomial in it has a degree d derivation, so $(\mathcal{P}, \mathcal{Q}, S)$ is d-complete.

To prove spectral-richness, we note that in [Fol01] the author gives an exact formula for each entry of the matrix $M = \int_S p(x)$ for any polynomial p. The formulas imply that $(n+d)! \pi^{-n/2} M$ is an integer matrix with entries (very loosely) bounded by $(n+d)! d! 2^n$. By Lemma 3.1, we conclude that S is δ-spectrally rich with $1/\delta = 2^{\text{poly}(n^d)}$.

We collect the examples discussed in this section here:

Corollary 3.5. *The following constraints admit rich solutions:*

- MAX-CSP*:* $\mathcal{P} = \{x_i^2 - x_i | i \in [n]\}$.
- MAX-CLIQUE*:* $\mathcal{P} = \{x_i^2 - x_i | i \in [n]\} \cup \{x_i x_j | (i,j) \notin E\}$.
- BALANCED SEPARATOR*:* $\mathcal{P} = \{x_i^2 - x_i | i \in [n]\}$, $\mathcal{Q} = \{2n/3 - \sum_i x_i, \sum_i x_i - n/3\}$.

- MATCHING: $\mathcal{P} = \{x_{ij}^2 - x_{ij} | i, j \in [n]\} \cup \{\sum_i x_{ij} - 1 | i \in [n]\} \cup \{x_{ij}x_{ik} | i, j, k \in [n]\}$.
- MAX-BISECTION: $\mathcal{P} = \{x_i^2 - x_i | i \in [n]\} \cup \{\sum_i x_i - n/2\}$.
- UNIT-VECTOR: $\mathcal{P} = \{\sum_i x_i^2 - 1\}$.

3.1 Limitations

While Theorem 4.1 allows us to prove that many different systems of polynomial constraints have well-behaved SoS proofs, there are a few areas where it comes up short. Most noticeably, to contain a rich set of solutions the solution space has to be nonempty. This can be a problem when trying to find SoS proofs of infeasibility. For example, one common technique is to introduce lower bounds on an objective function $f(x)$ of a maximization problem as constraints and attempt to use SoS to find a refutation, i.e. a proof of non-negativity for the constant polynomial -1. We are unable to show that these proofs can be taken to have polynomial bit complexity since they have empty solution spaces. As another example, we are unable to use our framework to show that refutations of the knapsack constraints use only polynomially many bits, even though it is clear by simply examining these known refutations that they only involve small coefficients.

4 Rich Solution Spaces Yield Bounded SoS Proofs

In this section we prove our main theorem:

Theorem 4.1. *Let $\mathcal{P} = \{p_1, \ldots, p_m\}$ and $\mathcal{Q} = \{q_1, \ldots, q_\ell\}$ be sets of polynomials with $S = \{\alpha \in \mathbb{R}^n | \forall p \in \mathcal{P} : p(\alpha) = 0, \forall q \in \mathcal{Q} : q(\alpha) \geq 0\}$. Assume that the set of solutions S is (k, δ, ε)-rich for $(\mathcal{P}, \mathcal{Q})$.*

Let $r(x)$ be a polynomial nonnegative on S, and assume r has a degree d sum-of-squares proof of nonnegativity

$$r(x) = \sum_{i=1}^{t_0} h_i^2 + \sum_{i=1}^{\ell} \left(\sum_{j=1}^{t_i} s_j^2 \right) q_i + \sum_{i=1}^{m} \lambda_i p_i.$$

Then r has a degree k sum-of-squares proof of nonnegativity such that the coefficients of every polynomial appearing in the proof are bounded by $2^{\text{poly}(n^k, \log \frac{1}{\delta}, \log \frac{1}{\varepsilon})}$. In particular, if S is rich then every coefficient can be written down with only $\text{poly}(n^d)$ bits.

Proof. First, we rewrite the proof into a more convenient form before proving bounds on each individual term. Because the elements of $\mathbf{v}$ are a basis for $\mathbb{R}[x]_d$, every polynomial in the proof can be expressed as $c^T\mathbf{v}$, where c is a vector of reals:

$$\begin{aligned} r(x) &= \sum_{i=1}^{t_0} (c_i^T \mathbf{v})^2 + \sum_{i=1}^{\ell} \left(\sum_{j=1}^{t_i} (d_{ij}^T \mathbf{v})^2 \right) q_i + \sum_{i=1}^{m} \lambda_i p_i \\ &= \langle C, \mathbf{v}\mathbf{v}^T \rangle + \sum_{i=1}^{\ell} \langle D_i, \mathbf{v}\mathbf{v}^T \rangle q_i + \sum_{i=1}^{m} \lambda_i p_i \end{aligned}$$

for PSD matrices C, $D_1, \ldots, D_\ell$. Next, we average this polynomial identity over all the points $\alpha \in S$:

$$E_{\alpha \in S}[r(\alpha)] = \langle C, E_{\alpha \in S}[\mathbf{v}(\alpha)\mathbf{v}(\alpha)^T] \rangle + \sum_{i=1}^{\ell} \langle D_i, E_{\alpha \in S}[\mathbf{v}(\alpha)\mathbf{v}(\alpha)^T] \rangle q_i(\alpha) + 0$$

The LHS is at most $\text{poly}(\|r\|, \|S\|)$, and the RHS is a sum of positive numbers, so the LHS is a bound on each term of the RHS. We would like to say that since S is δ-spectrally rich, the first term is at least $\delta Tr(C)$. Unfortunately the averaged matrix may have zero eigenvectors, and it is possible that C could have very large eigenvalues in these directions. However these eigenvectors must correspond to polynomials that are zero on S. Because $(\mathcal{P}, \mathcal{Q}, S)$ is complete, these can be absorbed into the final term. More formally, let $\Pi = \sum_u uu^T$ be the projector onto the zero eigenspace of $M = E_{\alpha \in S}[\mathbf{v}(\alpha)\mathbf{v}(\alpha)^T]$. Because $(\mathcal{P}, \mathcal{Q}, S)$ is complete, for each u we have a degree k derivation $u^T\mathbf{v} = \sum_i \sigma_{ui} p_i$. Then $\Pi \mathbf{v}\mathbf{v}^T = \sum_u (u^T\mathbf{v})u\mathbf{v}^T$. Thus we can write

$$\begin{aligned}
\langle C, \mathbf{v}\mathbf{v}^T \rangle &= \langle C, (\Pi + \Pi^\perp)\mathbf{v}\mathbf{v}^T(\Pi + \Pi^\perp)\rangle \\
&= \langle C, \Pi^\perp \mathbf{v}\mathbf{v}^T \Pi^\perp \rangle + \sum_u u^T\mathbf{v} \left(\langle C, \Pi^\perp \mathbf{v}u^T + \mathbf{v}u^T\Pi^\perp + \mathbf{v}u^T\Pi \rangle \right) \\
&= \langle \Pi^\perp C \Pi^\perp, \mathbf{v}\mathbf{v}^T \rangle + \sum_i \sigma_i p_i.
\end{aligned}$$

Doing the same for the other terms and setting $C' = \Pi^\perp C \Pi^\perp$ and similarly for D'_i, we get a new proof:

$$r(x) = \langle C', \mathbf{v}\mathbf{v}^T \rangle + \sum_{i=1}^{\ell} \langle D'_i, \mathbf{v}\mathbf{v}^T \rangle q_i + \sum_{i=1}^{m} \lambda'_i p_i.$$

Now after averaging over S, the zero eigenspace of C' is contained in the zero eigenspace of M. Taken with the δ-spectral richness, we have

$$\text{poly}(\|r\|, \|S\|) \geq \delta Tr(C) + \sum_{i=1}^{\ell} \delta Tr(D'_i) q_i(\alpha).$$

Because each $q_i(\alpha) \geq \varepsilon$, we get C' and D'_i have entries bounded by $\text{poly}(\|r\|, \|S\|, \frac{1}{\delta}, \frac{1}{\varepsilon})$.

The only thing left to do is to bound the coefficients λ'_i, but this is easy because the SoS proof is linear in these coefficients. If we imagine the coefficients of the λ'_i as variables, then the linear system induced by the polynomial identity

$$r(x) - \langle C', \mathbf{v}\mathbf{v}^T \rangle - \sum_{i=1}^{\ell} \langle D'_i, \mathbf{v}\mathbf{v}^T \rangle = \sum_{i=1}^{m} \lambda'_i p_i$$

is clearly feasible, and the coefficients of the LHS are bounded by $\text{poly}(\|r\|, \|S\|, \frac{1}{\delta}, \frac{1}{\varepsilon})$. There are $O(n^k)$ variables, so by Cramer's rule, the coefficients of the λ'_i can be taken to be bounded by $\text{poly}(\|\mathcal{P}\|^{n^k}, \frac{1}{\delta}, \frac{1}{\varepsilon}, \|r\|, \|S\|, n!)$. $\|\mathcal{P}\|, \|r\| \leq 2^{\text{poly}(n^d)}$ as they are considered part of the input, $\|S\| \leq 2^{\text{poly}(n^d)}$ by the explicitly bounded assumption, and $d \leq k$. Thus, this bound is at most $2^{\text{poly}(n^k, \log \frac{1}{\delta}, \log \frac{1}{\varepsilon})}$.

□

5 Boolean Systems With No Small-Coefficient Proofs

In [O17], the author gives an example of a polynomial system for which degree two SoS proofs can certify non-negativity of a certain poylnomial, but the proofs necessarily involves coefficients

of doubly-exponential size. However, there are two weaknesses in his example system. First, it is not a Boolean one, i.e. it contains variables y_i for which the constraint $y_i^2 - y_i = 0$ is not present in the constraints. Many practical optimization problems have Boolean constraints, and in [O17], the author hoped that having those constraints might suffice to imply that all proofs could have small bit complexity. Second, while the degree two proofs must have exponential bit complexity, there were degree four proofs of non-negativity with polynomial bit complexity. In this section, we strengthen his counterexample, giving an example of a Boolean system with n variables for which there is a polynomial that has a degree two proof of non-negativity, but no proof with polynomial bit complexity until degree $\Omega(\sqrt{n})$.

5.1 A First Example

The original example given in [O17] essentially contains the following system whose repeated squaring is responsible for the blowup of the coefficients in the proofs:

$$y_1^2 - y_2 = 0, \quad y_2^2 - y_3 = 0, \quad \ldots, \quad y_{n-1}^2 - y_n = 0, \quad y_n^2 = 0.$$

Clearly, the only solution to the system is $(0, 0, 0, \ldots, 0)$, and therefore the polynomial $\varepsilon - y_1$ must be non-negative over the solution space for any $\varepsilon > 0$. It is not as obvious whether or not an SoS proof of this non-negativity exists. It turns out that there is a degree two SoS proof as follows:

$$\begin{aligned}\varepsilon - y_1 \equiv \left(\sqrt{\frac{\varepsilon}{n}} - \left(\frac{n}{4\varepsilon}\right)^{1/2} y_1\right)^2 + \left(\sqrt{\frac{\varepsilon}{n}} - \left(\frac{n}{4\varepsilon}\right)^{3/2} y_2\right)^2 + \left(\sqrt{\frac{\varepsilon}{n}} - \left(\frac{n}{4\varepsilon}\right)^{7/2} y_3\right)^2 + \\ + \cdots + \left(\sqrt{\frac{\varepsilon}{n}} - \left(\frac{n}{4\varepsilon}\right)^{(2^n-1)/2} y_n\right)^2. \qquad (*)\end{aligned}$$

where the $\equiv$ is equality modulo the ideal generated by the constraints. Of course, this proof involves coefficients of doubly-exponential size, but one can prove that they are required. We will take $\varepsilon < 1/2$ for simplicity. We will define a linear functional $\phi : \mathbb{R}[Y]_d \to \mathbb{R}$ satisfying the following:

- $\phi[\varepsilon - y_1] = -\varepsilon$
- $\phi[p^2] \geq 0$ for any p^2 of degree at most d
- $\phi[\sigma_i(y_i^2 - y_{i+1})] = 0$ for any $i \leq n-1$ and σ_i of degree at most $d-2$
- $|\phi[\lambda y_n^2]| \leq (2\varepsilon)^{2^{n-1}} n^d \|\lambda\|$.

If such a ϕ exists, then for any degree d SoS proof of non-negativity

$$\varepsilon - y_1 = \sum_i h_i(y)^2 + \sum_{i=1}^{n-1} \sigma_i(y_i^2 - y_{i+1}) + \lambda \cdot y_n^2,$$

apply ϕ to both sides. We obtain $-\varepsilon \leq P + 0 + \phi[\lambda y_n^2]$, where $P \geq 0$. Because $|\phi[\lambda y_n^2]| \leq (2\varepsilon)^{2^{n-1}} n^d \|\lambda\|$, λ must contain a coefficient of size at least $\Omega(\frac{1}{n^d}\left(\frac{1}{2\varepsilon}\right)^{2^n})$.

To show that such a ϕ exists, we define it as follows. By the constraints, every monomial is equivalent to some power of y_1. For example, $y_1 y_2 y_3 \equiv y_1^7$. More generally, the constraints imply that $\prod_{i=1}^n y_i^{\beta_i} = y_1^{\sum_{j=1}^n 2^{j-1}\beta_j}$. Define ϕ by,

$$\phi\left(\prod_{i=1}^{n} y_i^{\beta_i}\right) = (2\varepsilon)^{\sum_i 2^{i-1}\beta_i}$$

One can easily check that this ϕ satisfies the above. Note that none of the variables y_i in the above system are boolean, which we achieve in the upcoming section.

5.2 A Boolean System

One simple way to try to make the system Boolean is to just add the constraints $y_i^2 = y_i$ to the system. Unfortunately, in that case it is easy to prove that $y_i - y_j = 0$ for each i and j, and of course $y_n = y_n^2 = 0$. It is too easy for SoS to figure out what each y_i should look like. Previously, the variables were unconstrained in any way, and we want to imitate that. We draw inspiration from the Knapsack problem, and we instead replace each instance of the variable y_i with a sum of $2k$ Boolean variables

$$y_i \to \sum_j w_{ij} - k,$$

and we consider the non-negative polynomial $\varepsilon - (\sum_j w_{1j} - k)$. Clearly there is a degree two proof of non-negativity for this polynomial since we can just replace each instance of y_i with $\sum_j w_{ij} - k$ in $(*)$.

It remains to show that there are no other proofs that have only small coefficients. Here, we use the fact that the Knapsack problem is hard for SoS: there is no SoS proof of degree less than $\Omega(k)$ that $\sum_j w_{ij} - k$ is not equal to any number $r \in (0, 1)$ [Gri01]. This allows us to use the Knapsack pseudodistribution to "pretend" that $\sum_j w_{ij} - k = (2\varepsilon)^{2^{i-1}}$. Specifically, for each $r \in (0, 1)$, there is a linear functional ϕ_r defined on polynomials of $2k$ Boolean variables which satisfies

- $\phi_r[\sigma_{ij}(w_{ij}^2 - w_{ij})] = 0$ for any σ_{ij} up to degree $O(k)$
- $\phi_r[\lambda \cdot ((\sum_j w_{ij} - k) - r)] = 0$ for any polynomial λ up to degree $O(k)$
- $\phi_r[p^2] \geq 0$ for any polynomial p^2 of degree at most $O(k)$.

Now, take the linear functional Φ defined on each polynomials of $2kn$ variables defined in the following way: Let $T = T_1 \cup T_2 \cup \cdots \cup T_n$ where T_i is a multiset that contains only the variables corresponding to y_i, and let w_T denote the associated monomial. Then define

$$\Phi[w_T] = \phi_{2\varepsilon}(w_{T_1})\phi_{(2\varepsilon)^2}(w_{T_2}) \dots \phi_{(2\varepsilon)^{2^{n-1}}}(w_{T_n}).$$

Clearly Φ is non-negative on squares and $\Phi[\sigma_{ij}(w_{ij}^2 - w_{ij})] = 0$ for any σ_{ij} up to degree $\Omega(k)$. Because $\Phi[\lambda(\sum_j w_{ij} - k)] = \Phi[(2\varepsilon)^{2^{i-1}}\lambda]$, Φ also satisfies $\Phi[\lambda((\sum_j w_{ij} - k)^2 - (\sum_j w_{i+1,j} - k))] = 0$ for each λ and $1 \leq i \leq n-1$. Finally, because each variable is Boolean, Φ of any monomial is at most one, so for any monomial w_M, $\Phi[w_M(\sum_j w_{nj} - k)^2] = \Phi[(2\varepsilon)^{2^{n-1}} w_M] \leq (2\varepsilon)^{2^{n-1}}$. There are at most $(nk)^d$ monomials, so $\Phi[\lambda(\sum_j w_{nj} - k)^2] \leq (nk)^d (2\varepsilon)^{2^{n-1}} \|\lambda\|$. Just as before, the existence of Φ implies that any degree d proof of non-negativity for $\varepsilon - (\sum_j w_{1j} - k)$ must contain coefficients of size at least $\Omega(\frac{1}{(nk)^d} \cdot \left(\frac{1}{2\varepsilon}\right)^{2^n})$. If we set $k = n$, then there are n^2 variables and no proof of non-negativity with coefficients smaller than doubly-exponential until degree n. This proves Theorem 1.2.

6 Max-Bisection Constraints

In this section, we prove our earlier claim that the Max-Bisection constraints admit rich solutions. Recall the constraints:

$$\mathcal{P}(n) = \left\{x_i^2 - x_i | i \in [2n]\right\} \cup \left\{\sum_i x_i - n\right\}.$$

Recall that to prove S is rich, we have to prove that it is spectrally rich, robust, and complete. Since the solution space lies in the hypercube, it is spectrally rich by Lemma 3.2, and it is clearly robust since $\mathcal{Q}$ is empty. It remains to prove that it is complete for some k. This proof follows a very similar path to [BBCH+16], due to the similar symmetry of the constraints.

Lemma 6.1. *$\mathcal{P}(n)$ is d-complete for any $d \leq n$.*

Proof. Let $S(n)$ denote the solution space of $\mathcal{P}(n)$, and let $M = \mathbb{E}_{\alpha \in S}[\mathbf{v}(\alpha)\mathbf{v}(\alpha)^T]$. Any zero eigenvector c of M can be associated with a polynomial $c^T\mathbf{v}$. Since

$$c^T M c = \sum_\alpha (c^T \mathbf{v}(\alpha))^2 = 0,$$

we must have $c^T v(\alpha) = 0$ for each $\alpha \in S$. We argue that any degree d polynomial which is identically zero on $S(n)$ must have a degree d derivation from $\mathcal{P}(n)$.

We proceed by induction on d. If $d = 0$, the only constant polynomial zero on $S(n)$ is the zero polynomial, which has the trivial derivation. Now consider the case of $d = c + 1$. We proceed in two parts. First, if r is fully symmetric, we show that it has a degree d derivation. Secondly, for any polynomial p which is zero on $S(n)$, we prove that $p - \frac{1}{(2n)!}\sum_{\sigma \in \mathcal{S}_n} \sigma p$ has a degree d derivation from $\mathcal{P}$, where σ acts on p by permuting the labels of the variables. Taken together, these two facts imply that r has a degree d derivation from $\mathcal{P}(n)$.

To prove the first part, note that a symmetric polynomial r is a linear combination of the elementary symmetric polynomials $e_1, \ldots, e_c$, and it is clear that $e_k(x)$ can be derived by taking the polynomial $(\sum_i x_i - n)^k$, reducing it to multilinear using the boolean constraints, and then reducing by $e_l(x)$ for each $l < k$. This will result in a constant polynomial, which must be the zero polynomial since we are only adding polynomials which are zero on $S(n)$, so the resulting polynomial must be zero on $S(n)$.

To prove the second part, let σ_{ij} be the transposition of labels i and j, and consider the polynomial $r - \sigma_{ij} r$. Writing $r = r_i x_i + r_j x_j + r_{ij} x_i x_j + q_{ij}$, where none of r_i,r_j,r_{ij}, nor q_{ij} depend on x_i or x_j, we can rewrite

$$r - \sigma_{ij} r = (r_i - r_j)(x_i - x_j).$$

Now because $r - \sigma_{ij} r$ evalutes to zero on any boolean string with exactly n ones, if we set $x_i = 1$ and $x_j = 0$, we know that $r_i - r_j$ is a polynomial that must evaluate to zero on any boolean string with exactly $n - 1$ ones. Because $\deg(r_i - r_j) = d - 1$, by the inductive hypothesis, $r_i - r_j$ has a degree $d - 1$ proof from $\mathcal{P}(n-1)$ (since $d \leq n$, clearly $d - 1 \leq n - 1$). This implies that $(r_i - r_j)(x_i - x_j)$ has a degree $d - 1$ proof from $\mathcal{P}(n)$:

$$\begin{aligned}(r_i - r_j)(x_i - x_j) &= \left[\sum_{t \neq i,j} \lambda_t \cdot (x_t^2 - x_t) + \lambda \cdot \left(\sum_{t \neq i,j} x_t - (n-1)\right)\right](x_i - x_j) \\ &= \sum_t \lambda'_t \cdot (x_t^2 - x_t) + \lambda \cdot \left(\sum_{t \neq i,j} x_t - (n-1) + (x_i + x_j - 1)\right)(x_i - x_j)\end{aligned}$$

$$= \sum_t \lambda'_t \cdot (x_t^2 - x_t) + \lambda' \cdot \left(\sum_t x_t - n \right)$$

where we used the fact that $(x_i + x_j - 1)(x_i - x_j) - (x_i^2 - x_i) + (x_j^2 - x_j) = 0$. The degree of this derivation is at most d because each λ_t has degree at most $d-3$, and $\lambda'_t = \lambda_t(x_i - x_j)$, and similarly for λ. Thus the inductive hypothesis implies that $r - \sigma_{ij} r$ has a degree d derivation, and since transpositions generate the symmetric group, this implies that $r - \frac{1}{(2n)!} \sum_{\sigma \in \mathcal{S}_n} \sigma r$ has a degree d proof from $\mathcal{P}(n)$. □

Remark 6.2. In this example, $\mathcal{P}$ is not a Gröbner basis for its ideal $\langle \mathcal{P} \rangle$. Indeed, the Gröbner basis for this ideal has exponential size. This is an example where our framework is applicable, even though Gröbner bases are intractable to compute.

References

[AL94] William Adams and Philippe Loustaunau, *An introduction to gröbner bases*, American Mathematical Society, 1994. 5

[BBCH+16] Gábor Braun, Jonah Brown-Cohen, Arefin Huq, Sebastian Pokutta, Prasad Raghavendra, Aurko Roy, Benjamin Weitz, and Daniel Zink, *The matching problem has no small symmetric sdp*, Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms (Philadelphia, PA, USA), SODA '16, Society for Industrial and Applied Mathematics, 2016, pp. 1067–1078. 6, 11

[BBH+12] Boaz Barak, Fernando G.S.L. Brandao, Aram W. Harrow, Jonathan Kelner, David Steurer, and Yuan Zhou, *Hypercontractivity, sum-of-squares proofs, and their applications*, Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC '12, ACM, 2012, pp. 307–326. 1

[BRS11] B. Barak, P. Raghavendra, and D. Steurer, *Rounding semidefinite programming hierarchies via global correlation*, Proc. FOCS, IEEE, 2011, pp. 472–481. 1

[BS14] Boaz Barak and David Steurer, *Sum-of-squares proofs and the quest toward optimal algorithms*, In Proceedings of the 2014 International Congress of Mathematicians. International Mathematical Union (2014). 1

[Fol01] Gerald B. Folland, *How to integrate a polynomial over a sphere*, The American Mathematical Monthly **108** (2001), no. 5, 446–448. 6

[Gri01] D. Grigoriev, *Complexity of positivstellensatz proofs for the knapsack*, computational complexity **10** (2001), no. 2, 139–154. 10

[GS11] Venkatesan Guruswami and Ali Kemal Sinop, *Lasserre hierarchy, higher eigenvalues, and approximation schemes for graph partitioning and quadratic integer programming with PSD objectives*, IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011, 2011, pp. 482–491. 1

[GV01] Dima Grigoriev and Nicolai Vorobjov, *Complexity of null-and positivstellensatz proofs*, Annals of Pure and Applied Logic **113** (2001), no. 1-3, 153–160. 1

[Las00] Jean Bernard Lasserre, *Optimisation globale et théorie des moments*, Comptes Rendus de l'Académie des Sciences-Series I-Mathematics **331** (2000), no. 11, 929–934. 1

[Las01] Jean B Lasserre, *Global optimization with polynomials and the problem of moments*, SIAM Journal on Optimization **11** (2001), no. 3, 796–817. 1

[Lau09] Monique Laurent, *Sums of squares, moment matrices and optimization over polynomials*, Emerging applications of algebraic geometry, Springer, 2009, pp. 157–270. 1

[Nes00] Yurii Nesterov, *Squared functional systems and optimization problems*, High performance optimization, Springer, 2000, pp. 405–440. 1

[O17] Ryan ODonnell, *Sos is not obviously automatizable, even approximately*, Innovations in Theoretical Computer Science (ITCS) (2017). 1, 3, 8, 9

[Par00] Pablo A Parrilo, *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*, Ph.D. thesis, California Institute of Technology, 2000. 1

[Sho87] Naum Z Shor, *Class of global minimum bounds of polynomial functions*, Cybernetics **23** (1987), no. 6, 731–734. 1